

THE Practical Engineer

Data visualization: the end of the rainbow

At the core of all good science and engineering is the respectful treatment of data: instruments are calibrated, algorithms are scrutinized, and the behavior of analytical or simulation models studied, on the assumption that the tools suit the data at hand.

But the knowledge to use one set of tools, data visualization software, is often lacking. Visualization should help make sense of the flood of output data. When applied without some insight into visual perception, however, it can introduce errors in understanding as surely as if a wrong analysis algorithm were used. In short, a picture can tell a thousand lies.

Blind spots in visualization are inevitable if attention is not paid to how the human visual system processes information, to the nature of the data, and to how the data are to be used. But once identified, these factors can be manip-

ulated to good effect, as a variety of visualizations drawn from fields as diverse as geography, medicine, and physics will show.

The key is the colormap, which may be defined as a mapping from a data value to a color. (The term colormap is also sometimes used in other contexts specifically for the contents of the memory locations governing display.)

A rainbow colormap is often supplied as a default in visualization software, the vendor wishing perhaps to provide the greatest possible range of colors with which to work. In this kind of colormap, red is mapped to the highest data value, blue to the lowest, and the other data values are interpolated along the full extent of the spectrum. An example would be a temperature profile mapped over a land mass on a weather map.

But there is more to color than meets the eye. Color, after all, is a perceptual as well as physical phenomenon. What is commonly called color—hue—is only one of three parameters. Another is the brightness of the signal—intensity. The third is the admixture of white—saturation. Change

BERNICE E. ROGOWITZ & LLOYD A. TREINISH
IBM Corp., Thomas J. Watson Research Center

any one parameter enough, and the color looks different. (The hue-intensity-saturation model of color is one of a several explored through the years, and captures some of the basic characteristics of basic color perception.) To make matters worse, the parameters' relationship to what is perceived is nonlinear. At the same intensity, for example, yellow appears brighter than blue.

Some of the perceptual principles involved have been implemented in software developed at IBM Corp.'s Thomas J. Watson Research Center, Yorktown Heights, N.Y. The module runs with IBM's visualization package Data Explorer and is called Pravda (for perceptual rule-based architecture for visualizing data accurately). Though not yet commercially available, it is in use in several projects [see "Using Pravda's rules," p. 59].

What happened to Florida?

Consider the pair of visualizations in Fig. 1. Both images represent the same data, but appear radically different—all the more surprising because they use colormaps that are mathematically equivalent, having a one-to-one mapping from color value to data point.

A rainbow colormap produces the left-hand view in Fig. 1. There, a large sheet of yellow dominates, an inlet of blue rimmed with cyan and green intrudes, and some dark red regions are clustered on the upper right. But treat the same data with another colormap, one perceptually tuned to the type of data and the message the visualization is to convey, and the contours of the southeastern United States leap into view [Fig. 1, right]. The boundary of the continental shelf and areas of deep ocean are shown in a purple that darkens with water depth. The Appalachian Mountains rising from the coastal plains are clearly distinguished as regions of lighter and lighter green.

Why the difference? The perceptually tuned colormap has been designed so that equal steps in the data variable will be perceived as equal steps in the representation. The data also include a threshold color value—a clear boundary of interest to the user of the data, namely, sea level. This characteristic of the interval data is also explicitly incorporated into the colormap.

In contrast, in the first image, the rainbow colormap has very rough contouring. Yet these blue, cyan, green,

yellow, and red regions purport to represent a fine sampling of a continuous phenomenon: the gradual rise of the land mass, and the gradual descent of the seabed (the sampling resolution is 1 meter vertically and about 8 km horizontally).

The rogues' gallery

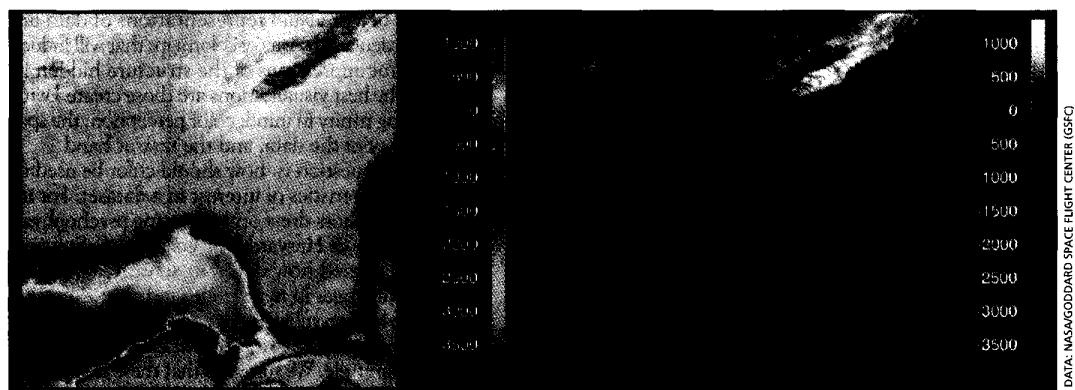
The misleading artifacts produced by the rainbow colormap are easy to recognize when a well-known continental coastline is the subject. With less familiar data, the visualization of features and relationships is trickier. A sense of how interpretations can go awry may be gleaned from the following rogues' gallery of visualizations taken from four disciplines: a portion of the Chesapeake Bay basin in Maryland; a magnetic resonance image (MRI) of a human head; noise from a jet aircraft engine; and the earth's magnetic field [Fig. 2].

Each row has three different visualizations of a single underlying dataset. The monochrome visualization on the left is a standard of comparison, as well as the simplest of the three images. In this type, a surface representation, each magnitude was placed at a height proportional to its value. Then perspective was added and the surface shaded with a light source at the location of the viewer. In short, the brightness levels, or luminances, are not analytically proportional to the datapoint values, but are a result of simple lighting.

The rogues' gallery comprises the centers of the rows; here the magnitudes are mapped, with the aid of a rainbow colormap, in a color value interpolated between blue and red. More appropriate visualizations using perceptual colormaps form the third image in each row, and will be discussed later, after a look at the rainbow mappings and their inadequacies.

The first portrait is an elevation model of a portion of the Chesapeake Bay basin in Maryland and part of Virginia, including the mouths of the Potomac and Patuxent rivers. With the surface representation [Fig. 2, top row, left] the viewer can see quite easily that the coastline emerges gradually and continuously, leading into the rivers' well-defined tributary structure.

The neighboring representation using the rainbow colormap masks many of these tributaries. Much of the data's



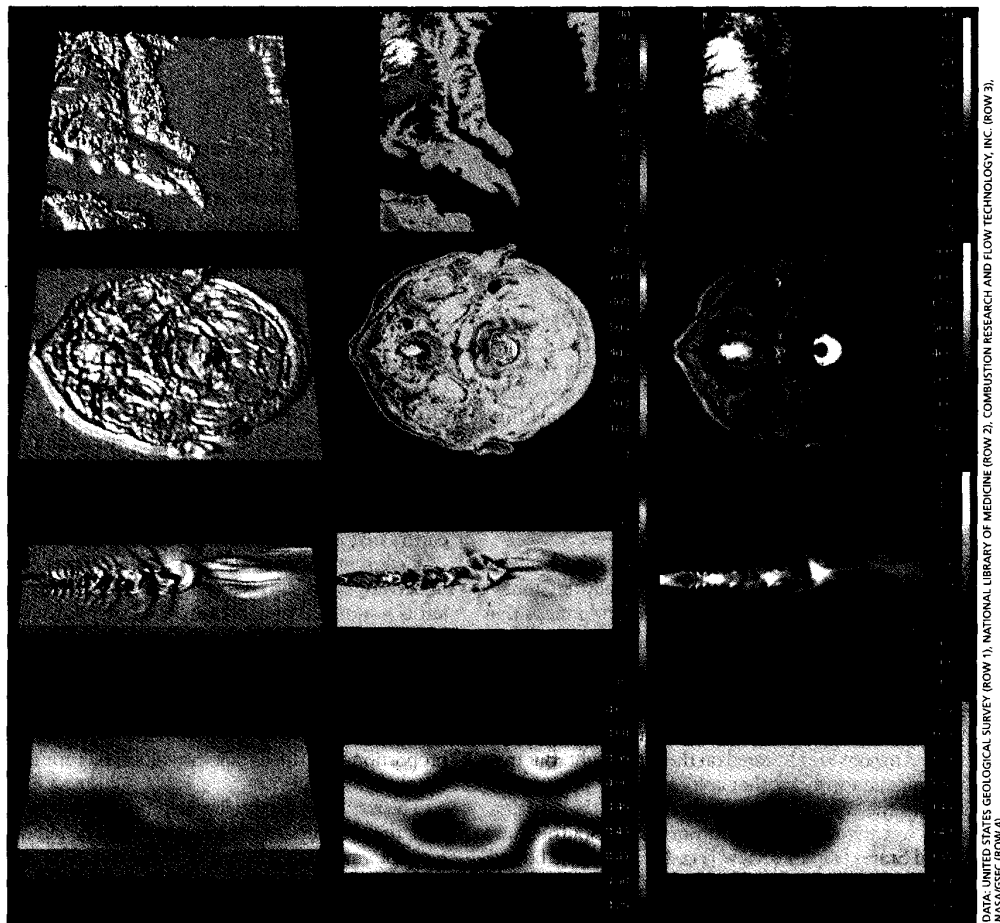
[1] These two images have the same underlying dataset of the Southeastern portion of the United States. The image at left, mapped to a rainbow spectrum common in visualization software, obscures the coastline and flattens the rise and fall of the land mass and sea bed. But in the right-hand image, a new color is chosen for the threshold at sea level, and the fine gradations are revealed by color changes matched to the appropriate mechanisms of visual perception.

[2] Each of these four rows contains three visualizations of the same data.

In the leftmost images, the magnitudes are raised to a height proportional to their value, producing the simplest visualizations.

The middle images, map the same data to rainbow colormaps, with loss of structural information.

For the rightmost visualizations, color-maps are chosen in accordance with the spatial frequencies of the data and how they are processed visually.



complexity is lost in the variation from green to cyan. The visualization also introduces a false segmentation of the higher elevations: instead of appearing continuous, they seem sharply divided into four distinct colored areas.

Data captured in an MRI scan of a human head appear next [Fig. 2, second row]. In the center image the entire numerical range from 150 to 300 appears to be a uniform cyan (the units refer to the intensity values of the scan image). Put another way, although these data change by a factor of two, all the values in this range look yellow. Similarly, all the values from 300 to 500 appear to be green.

What's more, the use of the rainbow colormap divides up the data range into just a few bands, suggesting incorrectly that there are just a few values in the data. To make matters worse, these bands are unevenly spaced. The colormap produces a contoured impression, masking the subtle, continuous variations in intensity easily visible in the monochrome surface representation.

A simulation of noise from a jet aircraft engine comes next. The leftmost representation [Fig. 2, third row] reveals a complex, turbulent scene. The viewer's understanding of the middle image, however, is led astray by the segmentation produced when the rainbow colormap is used. The image reveals neither the multiple peaks in the quickly changing visual detail of the region in the left half nor the undulating structure on the right.

Last comes a model of the earth's surface magnetic field on a cylindrical Cartesian projection [Fig. 2, bot-

tom row]. In this projection, the north pole is stretched along the top of the image, the south pole along the bottom, and the equator bisects the rectangle horizontally.

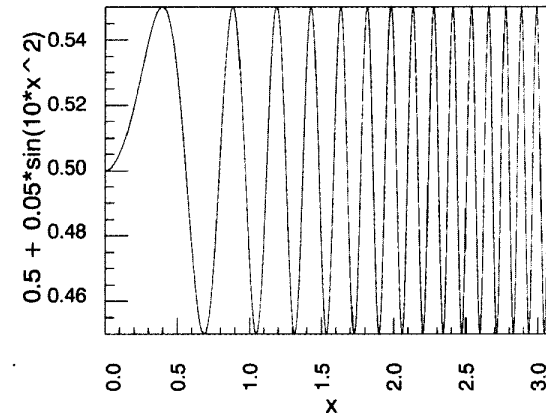
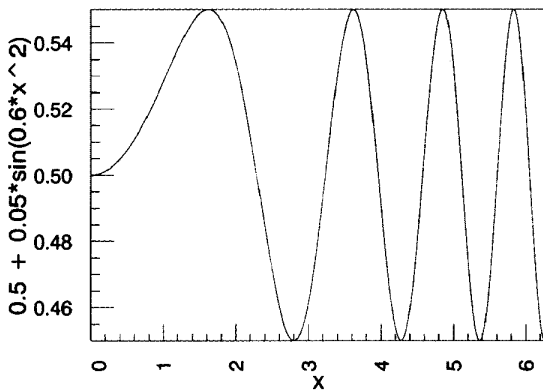
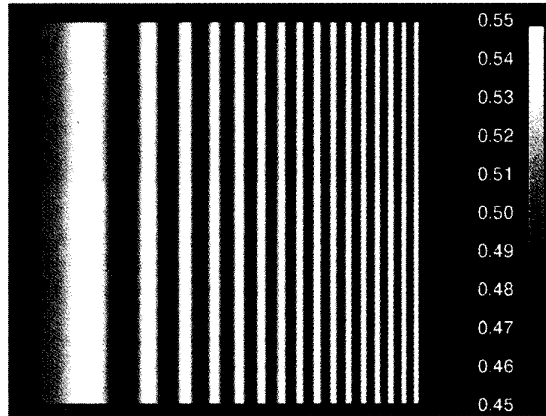
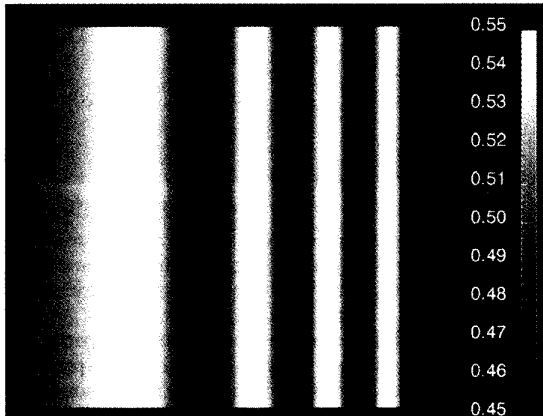
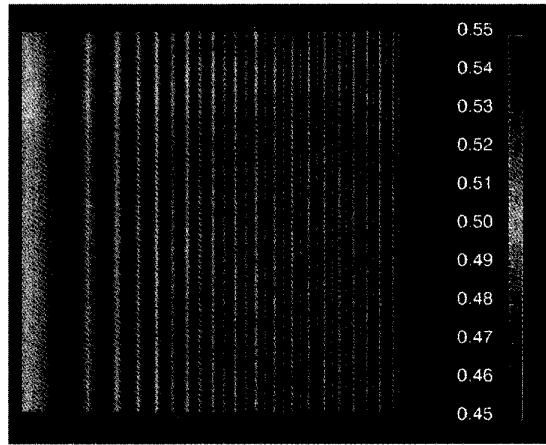
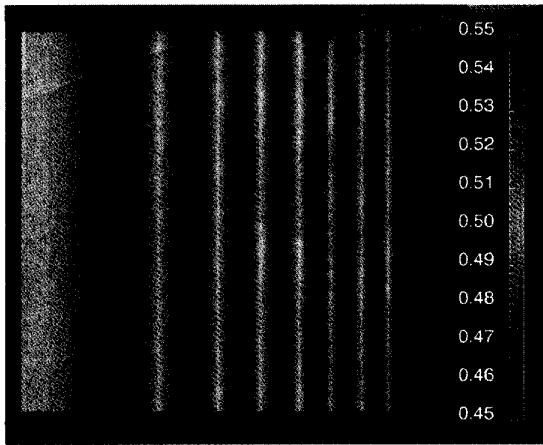
The surface representation shows how the magnetic field varies very gradually, increasing slightly at the magnetic poles. Yet the middle image strongly suggests large and sharp changes in magnetic field, particularly near the magnetic pole.

Perceptual encoding

Clearly, better visual representations of data are needed, particularly by way of colormaps that will induce more faithful impressions of the structure hidden in the data. The best visualizations are those created with the following trinity in mind: color perception, the spatial frequency of the data, and the task at hand.

The basic question is: how should color be used to encode characteristics of interest in a dataset. For an answer, the authors drew on pioneering psychophysical experiments of Harvard University's S. S. Stevens. Stevens examined how different sensory modes responded to changes in magnitude, and after extensive testing with human subjects, found that perceived luminance and perceived saturation varied smoothly with physical magnitude. We surmised that these dimensions would be useful for encoding magnitude in a colormap.

We extended these ideas to studying how colormaps based on these dimensions were employed depending on the spatial variation of the underlying data—that is, if the data has a high or low spatial frequency.



[3] Luminance and saturation mechanisms in human vision have a broad range in spatial sensitivity, but each has different strengths. Mapping the frequency-modulated grating [bottom row] only in saturation of two hues is better for perception of low frequency [top row]. The high-frequency cycles are more easily discerned with maps that vary in luminance [middle row].

Spatial vision

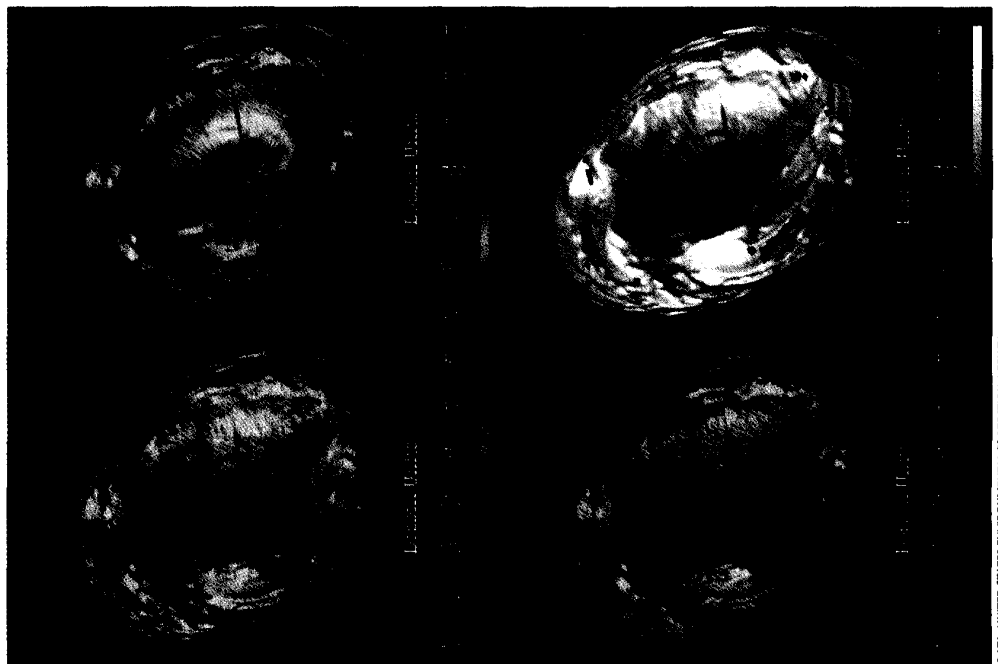
The luminance and saturation mechanisms in human vision each have a broad range in spatial sensitivity. But it turns out that the luminance mechanism excels at processing the detailed changes of high-frequency information, whereas saturation mechanisms are better for the undulating variations of low-frequency information [Fig. 3].

The bottom row of Fig. 3 is a waveform of a frequency-modulated grating, which begins at one cycle

and then increases. The variation in data value is represented by two colormaps: the first, shown on the top row, varies only in saturation and uses only two hues; the second, on the middle row, varies only in luminance (brightness).

At the low-frequency end of the spectrum [the left column of Fig. 4], the sinusoidal variation is more visible with the colormap that varies in saturation than with the colormap where luminance is the variable. Now consider the grating with the modulation of about

[4] Datasets like this atmospheric ozone distribution often have multiple spatial frequencies and can be examined with differing colormaps. Here, luminance changes reveal atmospheric circulation [upper right]; saturation changes in two directions reveal ozone depletion [lower left]; and the two maps can be combined [lower right]. A rainbow-mapped image is given for comparison [upper left].



10 cycles—a high spatial frequency. Here, many more of the grating's high-frequency cycles can be seen with the luminance-varying colormap than with the colormap that varies in saturation.

The use of perceptual colormaps with real data can now be examined. Consider the four images in Fig. 4, which have differing colormaps but the same dataset, taken from the atmosphere above the earth's southern hemisphere.

At top left, a rainbow colormap is shown for comparison. Next, the Dobson data (the units are used to measure ozone levels) is mapped onto a luminance-varying colormap, which reveals the data's high spatial frequencies quite successfully. Then, a saturation-varying colormap, representing low spatial-frequency data, was used to produce the third image; here the colormap increases in saturation from an achromatic (hue = 0) midpoint, becoming an ever more intense red for higher data values and an ever more intense green for lower data values. The grand finale combines the luminance and saturation variations of the low and high spatial-frequency maps.

More precisely, as can be seen in lower-right image, which uses both colormaps, ozone depletion, in general a low spatial frequency feature, is rendered with the saturation colormap. Atmospheric circulation, which is of high spatial frequency, is effectively captured by the luminance colormap.

The gallery revisited

The discussion of Fig. 2 was left hanging, with the four datasets represented only by a surface and a rainbow-colormapped visualization. Perceptual colormapping along the lines just described yields the last image in each row.

The Chesapeake topography very obviously supplies a threshold as well as interval data having a high spatial frequency. A perceptual colormap is therefore applied that uses luminance variation to reflect the rapid variation in data value across the domain. This variation increases from the threshold for the coastline

and decreases for elevations below a nominal sea level. Hue was used to reinforce the luminance variations. Notice how clearly this colormap shows the tributaries of the Potomac and Patuxent rivers.

In the MRI scan of the brain, the high spatial frequency of the interval data is enhanced with a primarily grayscale colormap. Luminance increases monotonically, and hue, which begins as a pure vivid blue, fades more and more to a pastel shade. This colormap produces a monotonic increase in perceived magnitude over the range. Notice how the fine detail in the structures in the image stands out in contrast with that lost with the rainbow colormap.

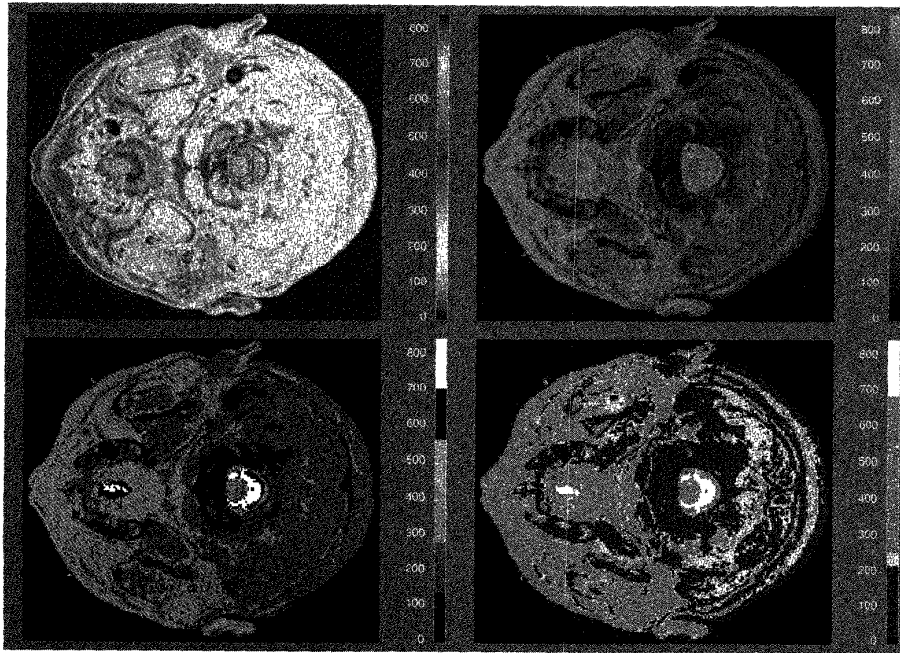
The two MRI images warrant a further word. The medical community has been cautious about adding color to its visualizations, and justifiably so. The rainbow-mapped visualization might appear to be prettier, but the (rarely provided) perceptually mapped one more accurately reflects the structure of the data.

A mix of data perceptions must be handled for the image of the simulated jet engine noise, because the data has regions of both low and high spatial frequency. A saturation-based colormap is used, but with a larger range than for the MRI task.

For the earth's magnetic field, a properly tuned colormap eliminates another kind of false impression that can be given by a rainbow colormap. Data with a low spatial frequency underlie this row. So a colormap was devised with a saturation increasing from achromatic (gray) in the yellow direction for higher magnetic field strength and in the blue direction for lower values. The magnetic field can then be seen to vary gradually—witness the steady increase in the range from 0.35 to 0.55 G—becoming strongest at the geomagnetic poles.

Tasks in visualization

In the examples covered so far, the interpretation depends on matching equal steps in the data to equal steps in perception—"equal" perceptually in the sense that one can say "x seems twice as dark as y," for example. These may be termed isomorphic visualization tasks.



[5] Ranges within a dataset are handled differently, depending on the visualization task. In isomorphic visualization, equal steps in the data are matched to equal perceptual steps [upper right]; segmentation strongly separates regions [lower right]; and highlighting draws attention to a specific region [lower right]. The rainbow-mapped image at the upper left is shown for comparison.

DATA: NATIONAL LIBRARY OF MEDICINE

It is also vital to match the visual representation to other visualization tasks: segmentation and highlighting. Both of these tasks reflect an interest in understanding how a phenomenon behaves within specific ranges of values. But within the total dataset, these ranges are used differently.

In segmentation, the analyst looks at the entire range of data, but the range is partitioned. The perceived ordering of the segments should match the order of the data values in some nonarbitrary way. Practically speaking, the analyst should get a sense of how data values and color values are related without having constantly to refer to the colormap.

In highlighting, the analyst identifies a limited range of a variable and sees how the range expresses itself in the dataset. The task may be to probe the exact ranges where the dose of a radiological treatment affects distant healthy tissue, or the particular magnitude at which the wind changes direction in a meteorological simulation.

Figure 5 exploits these refinements in a return to the skull cross-section of Fig. 2. Here, again the rainbow-mapped image [upper left] calls attention to the yellow areas not because they are in any way the most important, but because they are the brightest.

In contrast, the isomorphic color map at the upper right in Fig. 5 is designed to faithfully represent the data structure by matching equal steps in the data to equal perceptual steps. Notice that this colormap varies subtly from the isomorphic one in Fig. 2. In this case, the analyst is interested in the nut-shaped cluster near the center of the image.

In contrast to the image as a whole, this area has a fairly low spatial frequency structure. Since this data has mixed spatial frequency (as in reality do almost all datasets), but is leaning toward high, a luminance-based colormap was chosen, but with a relatively wide variation in hue and saturation.

Segmentation tasks are the job of the colormap at lower left, which must help delineate regions visually. Strongly separated and rather few color segments are

therefore called for, as seen in the color bar to the right.

Highlighting, handled by the colormap at lower right, draws the user's attention to certain regions in the image and their characteristic features. For example, perhaps the nut-shaped object toward the rear of the skull hides a tumor. The visualization in the upper right shows the same region, which is one of high spatial frequency. But to answer this particular question about this particular area, it is necessary to highlight the most important data values—those near the median of the range of interest—and the color map for the lower right image was designed accordingly, with a wide span straddling the median.

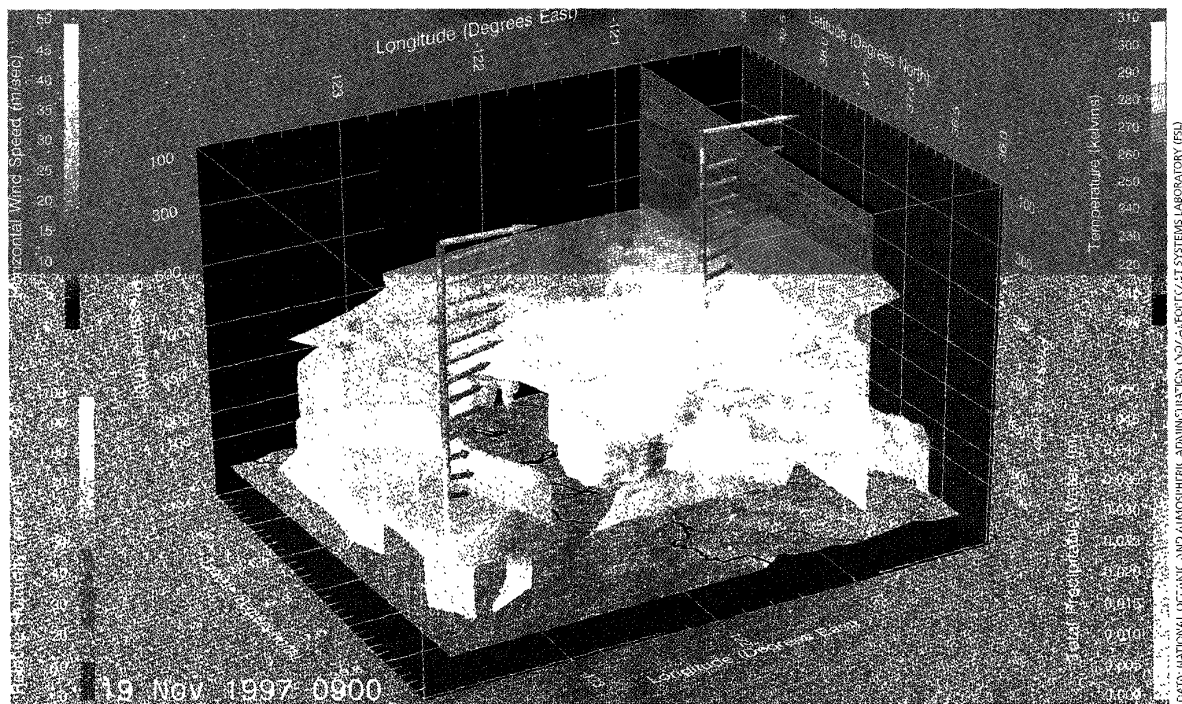
Three dimensions and up

The ideas discussed so far can be extended to visualizations in three dimensions with multiple datasets. An imposing example using assorted ways to clearly and logically present an abundance of information is shown in Fig. 6. These data are from an analysis of various weather observations on 19 November 1997 at 0900 local time in the San Jose, Calif., area. (This image, like many of the others shown, is one of a time series.) In this one image, four distinct colormaps are used to visualize temperature, precipitation, humidity, and wind speed.

These four variables, represented in some cases with a variety of geometries for multiple purposes, are joined by a variable for barometric pressure, the vertical axis in this information space. Atmospheric pressure is associated with elevation, and thus a topographic projection of the San Jose area can be situated with the addition of the other two axes, for latitude and longitude.

The choice of colormaps for each of these variables and their realizations is based in part upon their spatial characteristics and in part on the task associated with the visualization. For example, quite noisy (that is, high spatial frequency) data such as wind speed are mainly mapped into luminance. Quite smoothly varying data are mainly mapped into saturation, to impart a continuous representation.

For the task of segmenting the total humidity data



[6] This rich visualization of weather over San Jose, Calif., has differing colormaps to match multiple visualization tasks and different spatial frequencies of the data. Vector arrows representing the wind speed and direction are perched on poles of humidity readings; an isosurface of translucent white shows 90-percent-humidity levels ["clouds"]; and heavy lines mark coastlines [black] and rivers [blue].

levels over San Jose, the data are mapped into a set of bands (wrapped around the two "data poles") by way of the segmented colormap on the lower left of Fig. 6, which maintains the perception of ordinal relationships in the data.

Another segmentation approach to the relative-humidity data is in the isosurface visualization in white of the values of 90 percent; this surface is translucent so as not to obscure other information. The white surfaces, logically enough considering the rough conditions under which clouds are formed, are easily recognized as cloud boundaries. The temperatures at all heights at a given longitude are shown as segments on a vertical slice: temperature is smoothly varying data, but this task is segmentation, and thus requires more segments.

Vector arrows (or, in visualization terms, glyphs) representing the upper air wind data are posted along the two poles that also serve for the relative-humidity data. These point in the direction of the wind field and their color is correlated with the horizontal wind speed. The speed values, represented with a purplish wind-speed colormap [upper left], are also represented by the varying lengths of the arrows.

The relation of the full set of relative humidity data and the 90-percent-mark isosurface making up the "clouds" subset can be seen where each pole intersects the contour: a check of the humidity colormap [lower left] shows the color at this point on the bar and pole as 90. The pole-and-arrow combination appears twice, indicative of readings at different locations.

A variable for total precipitation [see the cyan-pink colormap at lower right] is overlaid on the topographic map at the bottom of the image. The precipitation data

is low spatial frequency, and thus the hues are altered primarily by saturation.

To show in general how the precipitation is deposited, the median of the values (0.025) was assigned the maximum desaturation, with the red and blue hues becoming more saturated as they increase and decrease, respectively. Atop the topographic map run the paths of rivers [blue] and coastlines [black].

In sum, modern systems for creating visualizations have evolved tremendously. But their use is very much an improvisation, with analysts fiddling with them repeatedly until something satisfactory seems in hand. A better grasp of just what is involved can improve one's effectiveness in creating visualizations, and, more importantly, the viewer's understanding of them. ♦

To probe further

The classic article describing the notion of perceptual scales, and how different stimuli can be used to represent magnitude information, is S. S. Stevens' "Matching Functions Between Loudness and Ten Other Continua." Anthologized frequently, it first appeared in *Perception and Psychophysics*, Vol. 1, pp. 5-8, 1966.

An early paper describing the importance of colormaps in computers is P. K. Robertson's "Visualizing Color Gamuts: A User Interface for the Effective Use of Perceptual Color Spaces in Data Displays," *IEEE Computer Graphics and Applications*, Vol. 8, pp. 50-63, September 1988.

The importance of selecting colormaps that allow the user to see differences in data over a large dynamic range is described in the same journal by H. Lefkowitz and G. T. Herman, in "Color Scales for Image Data" *IEEE Computer Graphics and Applications*, Vol. 12, no. 1, pp. 72-80, January 1992.

Using Pravda's rules

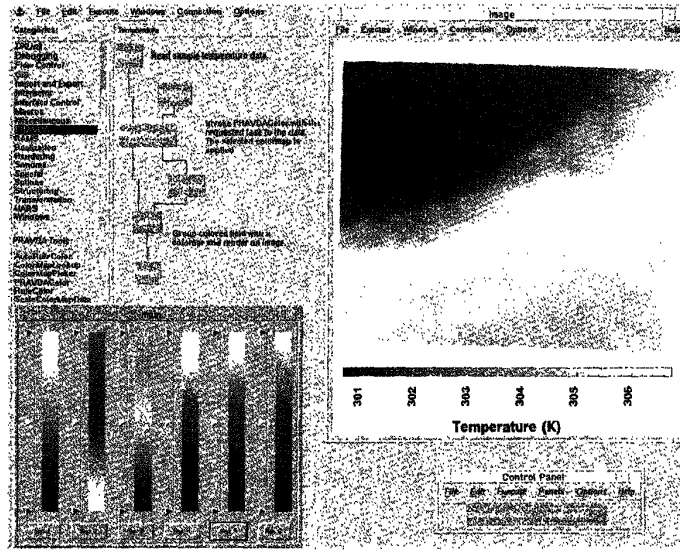
Few users of visualization tools want to become experts in human perception. In software called Pravda (perceptual rule-based architecture for visualizing data accurately), choices for colormapping are offered to the user based on principles of perception and color theory.

Shown here is Pravda being used during an interactive session with IBM Visualization Data Explorer (DX) software. DX is a visual-language, object-based package, as can be recognized from the connected modules shown in green.

The data to be visualized are imported into Data Explorer and flow into a module called PravdaColor. This tool determines the data's characteristics, including their spatial frequency—essentially, checking the results of data filtering.

With the aid of a control panel [lower right], the user selects the colormapping goal of the visualization: isomorphic, segmentation, or highlighting. These choices call up rules that constrain the set of colormaps. Here, six colormaps have been offered.

The data are from a temperature model for weather. The user has selected an isomorphic task (equal intervals represented by equal steps in perceived color).



And, since the data are of a low spatial frequency, these colormaps all encode variations in magnitude mostly as saturation variations. A click on any of the colormaps applies it directly to the data. The user is free to vary the range of the colormap, and here the full range of the fifth colormap has been selected.

Pravda is not available commercially, but a version of it may be in the future.

—B.E.R. & L.A.T.

The authors' "Using Perceptual Rules in Interactive Visualization" is a more comprehensive look into the psychophysical principles described in this article. The paper appeared in the *SPIE Proceedings on Human Vision, Visual Processing and Digital Display*, Vol. 2179, 1994, pp. 287–95.

Among larger texts, E. Tufte's *The Visual Display of Quantitative Information* (Graphics Press, Cheshire, Conn. 1990) has quickly become a classic, with beautiful and superbly chosen examples plucked from this field's long history.

As for standard techniques and methods in computer graphics, including treatment of color spaces, no library is complete without J. D. Foley, A. van Dam, S. K. Feiner, and S. F. Hughes's *Computer Graphics, Principles and Practice*, second edition (Addison-Wesley, Reading, Mass., 1991).

S.S. Steven's *Handbook of Experimental Psychology* (John Wiley & Sons, New York, 1955) has an excellent introductory chapter on measurement theory and data types. For a modern treatment of vision from a psychophysical perspective, see Brian A. Wandell's *Foundations of Vision* (Sinauer Associates, Sunderland, Mass., 1995).

The authors' implementation of Pravda (perceptual rule-based architecture for visualizing data accurately) is discussed in a paper by them and L. Bergman in "A Rule-based Tool for Assisting Colormap Selection" *Proceedings of the IEEE Computer Society Visualization '95*, October 1995, pp. 118–125, and can also

be viewed on the World Wide Web at <http://www.almaden.ibm.com/dx/vis96/proceedings/PRAVDA/index.htm>.

Pravda was designed for IBM Visualization Data Explorer (DX), a general-purpose software package for interactive data visualization and analysis. DX has a graphical program editor that allows the user to create a visualization using a point and click interface. DX runs on seven major Unix platforms and on Intel-based PCs, and is parallelized for multiprocessor systems. More information about Data Explorer is available on the Web at www.ibm.com/dx.

About the authors

Bernice E. Rogowitz (M) manages the visual analysis group at IBM Corp.'s Thomas J. Watson Research Center, in Yorktown Heights, N.Y. She is also founder and co-organizer of the Conference on Human Vision and Electronic Imaging within SPIE—The International Society for Optical Engineering/Imaging Science & Technology.

Lloyd A. Treinish (M) is a research staff member in the same IBM group at the Thomas J. Watson Research Center. His focus is on data visualization and methods of data management for earth, environmental, and space sciences. He is editor of the "Visualization Blackboard" column in *IEEE Computer Graphics and Applications*, and was program co-chair of the IEEE Visualization '98 conference.

Spectrum editor: Robert Braham